

---

# Quantized neural network design under weight capacity constraint

---

Sungho Shin, Kyuyeon Hwang, and Wonyong Sung

Department of Electrical and Computer Engineering  
Seoul National University

Seoul, 08826 Korea

sungho.develop@gmail.com, kyuyeon.hwang@gmail.com, wysung@snu.ac.kr

## Abstract

The complexity of deep neural network algorithms for hardware implementation can be lowered either by scaling the number of units or reducing the word-length of weights. Both approaches, however, can accompany the performance degradation although many types of research are conducted to relieve this problem. Thus, it is an important question which one, between the network size scaling and the weight quantization, is more effective for hardware optimization. For this study, the performances of fully-connected deep neural networks (FCDNNs) and convolutional neural networks (CNNs) are evaluated while changing the network complexity and the word-length of weights. Based on these experiments, we present the *effective compression ratio* (ECR) to guide the trade-off between the network size and the precision of weights when the hardware resource is limited.

## 1 Introduction

Deep neural networks (DNNs) begin to find many real-time applications, such as speech recognition, autonomous driving, gesture recognition, and robotic control (Sak et al., 2015; Chen et al., 2015; Shin & Sung, 2016; Corradini et al., 2015). Recent works show that the precision required for implementing fully-connected deep neural networks (FCDNNs), convolutional neural networks (CNNs) or recurrent neural networks (RNNs) needs not be very high, especially when the quantized networks are trained again to learn the effects of lowered precision. In the fixed-point optimization examples shown in Hwang & Sung (2014), neural networks with ternary weights showed quite good performance which was close to that of floating-point arithmetic. However, the performance of DNNs usually degrades when the weights are represented using a very low precision. Thus, we have a question whether it might be a better option to reduce the network size, instead of severely quantizing the weights, for efficient implementations.

In this work, we compare the performance of FCDNNs and CNNs under two constraints for hardware implementation, one is reducing the network size and the other is lowering the precision of the weights. We conduct the experiments with FCDNNs for phoneme recognition and CNNs for image classification. To control the network complexity, the number of units in each hidden layer is varied in the FCDNN. For the CNN, the number of feature maps for each layer is changed. The retraining based quantization algorithm is used for fixed-point optimization of weights (Hwang & Sung, 2014).

Based on the experiments, we propose a metric called the *effective compression ratio* (ECR) that compares the complexity of floating-point and fixed-point networks showing the same performance. This analysis intends to provide a guideline to network size and word-length determination for efficient hardware implementation of deep neural networks (DNN).

## 2 Related Work

Fixed-point design of DNNs with ternary weights show quite good performances that are very close to the floating-point results (Hwang & Sung, 2014; Anwar et al., 2015a; Shin et al., 2016). The ternary weight based FCDNN is used for VLSI implementations, by which the algorithms can operate with only on-chip memory consuming very low power (Kim et al., 2014). The CNN is implemented by XNOR-bitcounting operations (Rastegari et al., 2016). Binary weight based deep neural network design is also studied (Courbariaux et al., 2015). Pruned floating-point weights are utilized for efficient GPU-based implementations, where small valued or less important weights are forced to zero to reduce the number of arithmetic operations and the memory space for weight storage (Yu et al., 2012; Han et al., 2015; Anwar et al., 2015b).

Most of the above works are experimented using large size neural networks. However, mobile or embedded portable devices have limited resources, and thus small size fixed-point networks showing good performances are very needed.

## 3 Fixed-Point FCDNN and CNN Design

This section explains the design of FCDNN and CNN with varying network complexity and weight precision.

### 3.1 FCDNN and CNN Design

In this work, we examine an FCDNN for phoneme recognition and a CNN for image classification. The reference DNN has four hidden layers. Each of the hidden layers has  $N_h$  units; the value of  $N_h$  is altered to control the complexity of the network. We conduct experiments with  $N_h$  value of 32, 64, 128, 256, 512, and 1024. The input layer of the network has 1,353 units to accept 11 frames of a Fourier-transform-based filter-bank with 40 coefficients (+energy) distributed on a mel-scale, together with their first and second temporal derivatives. The output layer consists of 61 softmax units which correspond to 61 target phoneme labels (Mohamed et al., 2012). Phoneme recognition experiments were performed on the TIMIT corpus.

The CNN used is for CIFAR-10 dataset (Krizhevsky & Hinton, 2009). It contains a training set of 50,000 and a test set of 10,000 images. We divided the training set to 40,000 images for training and 10,000 images for validation. The reference CNN has 3 convolution and max-pooling layers, a fully connected hidden layer with 64 units, and the output with 10 softmax units. We control the number of feature maps in each convolution layer. The reference size has 32-32-64 feature maps with a 5 by 5 kernel size as used in Krizhevsky (2014). To know the effects of network size variation, the number of feature maps is reduced or increased. The configurations of the feature maps used for the experiments are 8-8-16, 16-16-32, 32-32-64, 64-64-128, 96-96-192, and 128-128-256. Note that the fully connected layer in the CNN is not changed.

### 3.2 Fixed-Point Optimization of DNNs

Reducing the word-length of weights brings several advantages in hardware based implementation of neural networks. First, it lowers the arithmetic precision, and thereby reduces the number of gates needed for multipliers. Second, the size of memory for storing weights is minimized, which would be a big advantage when keeping them on a chip, instead of external DRAM or NAND flash memory. Note that FCDNNs demand a very large number of weights. Third, reduced arithmetic precision or minimization of off-chip memory accesses leads to low power consumption.

The fixed-point DNN algorithm design consists of three steps: floating-point training, direct quantization, and retraining of weights. Refer to Hwang & Sung (2014) for the details.

## 4 Analysis of Quantization Effects

The fixed-point performance of the FCDNN is shown in Figure 1a, where the number of hidden units in each layer varies. The performance of direct 2 bits (ternary levels), direct 3 bits (7-levels), retrain-based 2 bits, and retrain-based 3 bits are compared with the floating-point results. We can

find that the performance gap between the floating-point and the retrain-based fixed-point networks converges very fast as the network size grows. Direct quantization does not show good results at any network size. In this figure, the performance of the floating-point network almost saturates when the network size is about 1024. Note that the TIMIT corpus that is used for training has only 3 hours of data. Thus, the network with 1024 hidden units can be considered in the ‘*training-data limited region*’. Here, the gap between the floating-point and the fixed-point networks almost vanishes when the network is in the ‘*training-data limited region*’. However, when the network size is limited, such as 32, 64, 128, or 256, there exists some performance gap between the floating-point and highly quantized networks even if retraining on the quantized networks is performed. The similar

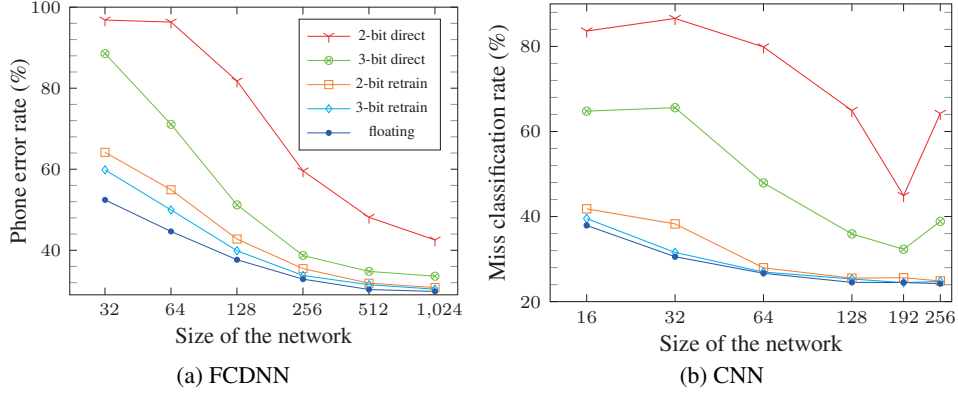


Figure 1: Comparison of retrain-based and direct quantization. All weights are quantized with ternary and 7-level weights. Note that each FCDNN has four hidden layers. In the figure (b), x-axis label “16, 32, 64, 128, 192, 256” represents the number of feature map is “8-8-16, 16-16-32, 32-32-64, 64-64-128, 96-96-192, 128-128-256”.

experiments are conducted for the CNN with varying feature map sizes, and the results are shown in Figure 1b. The configurations of the feature maps used for the experiments are 8-8-16, 16-16-32, 32-32-64, 64-64-128, 96-96-192, and 128-128-256. The size of the fully connected layer is not changed. In this figure, the floating-point and the fixed-point performances also converge very fast as the number of feature maps increases. The floating-point performance saturates when the feature map size is 128-128-256, and the gap is less than 1% when comparing the floating-point and the retrain-based 2-bit networks. This suggests that a fairly high-performance feature extraction can be designed even using very low-precision weights if the number of feature maps can be increased.

## 5 Efficient DNN Design with Hardware Constraints

As the number of quantization levels decreases, the memory space needed is reduced at the cost of sacrificing the accuracy. Therefore, there can be a trade-off between the network size reduction and aggressive quantization. Figure 2a shows the framewise phoneme error rate on TIMIT corpus while varying the layer size of FCDNNs with a various number of quantization bits from 2 to 8 bits. Note that the network has four hidden layers containing the same number of units.

In this section, we propose a guideline for finding the optimal bit-widths when the desired accuracy or the network size is given. Note that we assume  $2^n - 1$  quantization levels are represented by  $n$  bits (i.e. 2 bits are required for representing a ternary weight). For simplicity, all layers are quantized with the same number of quantization levels. Based on this observation, we introduce a metric called the *effective compression ratio (ECR)*, which is defined as follows:

$$ECR = \left[ \frac{\text{Size of a floating-point network}}{\text{Size of a quantized network showing the same performance}} \right] \quad (1)$$

Figure 2b describes how to compare the hardware efficiency of floating-point and fixed-point networks. In this figure, we assume the target performance of 32.87% which can be obtained using a floating-point DNN with the network size of 256. This graph shows that the target performance can be obtained with the network size of  $N1$  when 3 bits weights are used, and that of  $N2$  when 2 bits

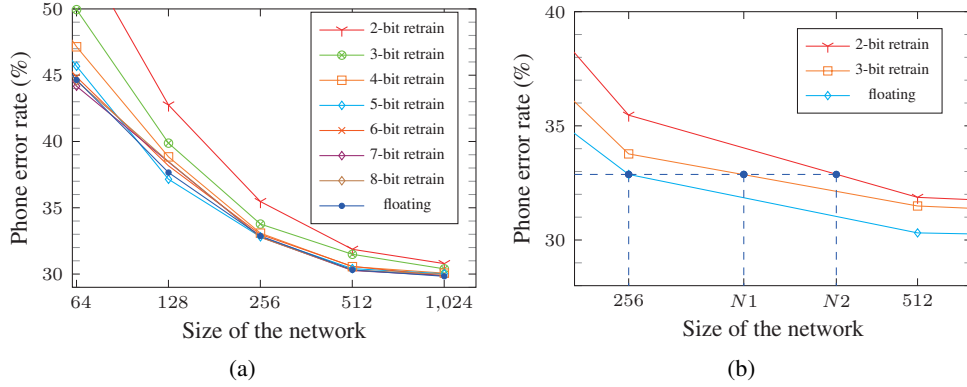


Figure 2: (a) Framewise phone error rate of phoneme recognition FCDNNs with respect to the size of the networks for weights with quantization. (b) Obtaining an effective number of parameters for the uncompressed network.

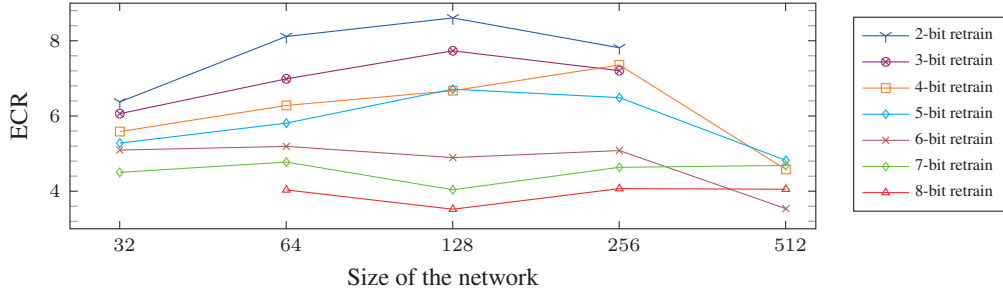


Figure 3: Effective compression ratio (ECR) of fixed-point networks with respect to the layer size of reference floating-point networks.

quantization is employed. Thus, the compression obtained by 3-bit quantization can be roughly figured as  $(32 \times 256^2)/(3 \times N1^2)$ . If  $N1$  is very close to 256, there can be about 10  $(= 32/3)$  times compression. But, if  $N1$  is 512, the compression drops to only about 2.5 because the number of parameters of FCDNN is proportional to the square of the network size.

The ECRs for various network sizes and quantization bits are shown in Figure 3. The figure illustrates that the 2-bit or (maybe 1-bit) quantization can lead to the best compression ratio when the target performance is low, which means a high phoneme error rate in this example. However, for designing a fairly high-performance network, increasing the network size with severe quantization does not yield hardware efficient networks. The optimum number of bits for obtaining the performance corresponding to that of 512 sizes floating-point DNNs is 4 or 5 bits. Further reducing the word-length demands very increased size networks, and as a result, the total number of bits increases.

## 6 Conclusion

Hardware efficient deep neural networks can be designed either by lowering the number of units in each layer or reducing the number of bits for weight quantization. We evaluate the performance of fixed-point deep neural networks and analyze the trade-off between the complexity and the precision of the weights. This study shows that low-performance hardware efficient DNNs can be designed with severely quantized weights. In the low-performance region, the DNN performance increases very rapidly as the network size grows. Thus, it is possible to compensate for the quantization effects by slightly increasing the network size. However, for a high-performance DNN design, compensation of quantization effects by increasing the network size is difficult, and thus severe quantization does not lead to efficient hardware design. The effective compression ratio is given for a DNN design when the network size and the precision vary.

## Acknowledgments

This work was supported in part by the Brain Korea 21 Plus Project and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (No. 2015R1A2A1A10056051).

## References

- Anwar, Sajid, Hwang, Kyuyeon, and Sung, Wonyong. Fixed point optimization of deep convolutional neural networks for object recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pp. 1131–1135. IEEE, 2015a.
- Anwar, Sajid, Hwang, Kyuyeon, and Sung, Wonyong. Structured pruning of deep convolutional neural networks. *arXiv preprint arXiv:1512.08571*, 2015b.
- Chen, Chenyi, Seff, Ari, Kornhauser, Alain, and Xiao, Jianxiong. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2722–2730, 2015.
- Corradini, Maria Letizia, Giantomassi, Andrea, Ippoliti, Gianluca, Longhi, Sauro, and Orlando, Giuseppe. Robust control of robot arms via quasi sliding modes and neural networks. In *Advances and Applications in Sliding Mode Control systems*, pp. 79–105. Springer, 2015.
- Courbariaux, Matthieu, Bengio, Yoshua, and David, Jean-Pierre. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in Neural Information Processing Systems*, pp. 3105–3113, 2015.
- Han, Song, Mao, Huizi, and Dally, William J. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. *CoRR*, abs/1510.00149, 2, 2015.
- Hwang, Kyuyeon and Sung, Wonyong. Fixed-point feedforward deep neural network design using weights +1, 0, and -1. In *Signal Processing Systems (SiPS), 2014 IEEE Workshop on*, pp. 1–6. IEEE, 2014.
- Kim, Jonghong, Hwang, Kyuyeon, and Sung, Wonyong. X1000 real-time phoneme recognition VLSI using feed-forward deep neural networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pp. 7510–7514. IEEE, 2014.
- Krizhevsky, A. CUDA-convnet, 2014.
- Krizhevsky, Alex and Hinton, Geoffrey. Learning multiple layers of features from tiny images. 2009.
- Mohamed, Abdel-rahman, Dahl, George E, and Hinton, Geoffrey. Acoustic modeling using deep belief networks. *Audio, Speech, and Language Processing, IEEE Transactions on*, 20(1):14–22, 2012.
- Rastegari, Mohammad, Ordonez, Vicente, Redmon, Joseph, and Farhadi, Ali. XNOR-Net: Imagenet classification using binary convolutional neural networks. *arXiv preprint arXiv:1603.05279*, 2016.
- Sak, Haşim, Senior, Andrew, Rao, Kanishka, and Beaufays, Françoise. Fast and accurate recurrent neural network acoustic models for speech recognition. *arXiv preprint arXiv:1507.06947*, 2015.
- Shin, Sungho and Sung, Wonyong. Dynamic hand gesture recognition for wearable devices with low complexity recurrent neural networks. In *Circuits and Systems (ISCAS), 2016 IEEE International Symposium on*, pp. 2274–2277. IEEE, 2016.
- Shin, Sungho, Hwang, Kyuyeon, and Sung, Wonyong. Fixed-point performance analysis of recurrent neural networks. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 976–980. IEEE, 2016.
- Yu, Dong, Seide, Frank, Li, Gang, and Deng, Li. Exploiting sparseness in deep neural networks for large vocabulary speech recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pp. 4409–4412. IEEE, 2012.